

08-18-00

A

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE****CERTIFICATE OF EXPRESS MAILING**

I hereby certify that this paper and the documents and/or fees referred to as attached therein are being deposited with the United States Postal Service on August 16, 2000 in an envelope as "Express Mail Post Office to Addressee" service under 37 CFR §1.10, Mailing Label Number **EL631003771US**, addressed to the Assistant Commissioner for Patents, Washington, DC 20231.

Attorney Docket No.:

SUN1P275/P4783

First Named Inventor:

David Wallman

Mandy Franklin  
Mandy Franklin

**UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR. § 1.53(b))**

Assistant Commissioner for Patents  
Box Patent Application  
Washington, DC 20231

☐ Duplicate for  
fee processing

**Re:** This is a request for filing a patent application under 37 CFR. § 1.53(b) in the name of inventors:  
David Wallman

**For:** METHOD AND APPARATUS FOR CACHING NATIVE CODE IN A VIRTUAL MACHINE  
INTERPRETER

Application Elements:

- ☒ 26 Pages of Specification, Claims and Abstract  
☒ 11 Sheets of **formal** Drawings  
☒ 02 Pages Combined Declaration and Power of Attorney

Accompanying Application Parts:

- ☒ Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)  
☐ 37 CFR 3.73(b) Statement by Assignee  
☐ Information Disclosure Statement with Form PTO-1449  
☐ Copies of IDS Citations  
☐ Preliminary Amendment  
☒ Return Receipt Postcard  
☐ Small Entity Statement(s)  
☐ Other:

Fee Calculation (37 CFR § 1.16)

	(Col. 1) <u>NO. FILED</u>	(Col. 2) <u>NO. EXTRA</u>	<u>SMALL ENTITY</u> <u>RATE FEE</u>	<u>OR</u>	<u>LARGE ENTITY</u> <u>RATE FEE</u>
BASIC FEE			\$345 \$	OR	\$690 \$690
TOTAL CLAIMS	<u>24</u> -20 = <u>04</u>		x 9 = \$	OR	x 18 = \$ 72
INDEP CLAIMS	<u>05</u> -03 = <u>02</u>		x 39 = \$	OR	x 78 = \$156
[ ] Multiple Dependent Claim Presented			\$130 = \$	OR	\$260 = \$
* If the difference in Col. 1 is less than zero, enter "0" in Col. 2.			Total \$	OR	Total \$918

☒ Check No. 8682 in the amount of \$958.00 is enclosed.

☒ The Commissioner is authorized to charge any fees beyond the amount enclosed which may be required, or to credit any overpayment, to Deposit Account No. 50-0388 (Order No. SUN1P275).

General Authorization for Petition for Extension of Time (37 CFR §1.136)

☒ Applicants hereby make and generally authorize any Petitions for Extensions of Time as may be needed for any subsequent filings. The Commissioner is also authorized to charge any extension fees under 37 CFR §1.17 as may be needed to Deposit Account No. 50-0388 (Order No. SUN1P275).

☒ Please send correspondence to the following address:

**Customer Number 022434**  
**BEYER WEAVER & THOMAS, LLP**  
 P.O. Box 130  
 Mountain View, CA 94042-0130  
 Telephone (650) 961-8300  
 Fax (650) 961-8301



Date: August 16, 2000

Peggy A. Su  
 Peggy A. Su  
 Registration No. 41,336

APPLICATION FOR U.S. PATENT

METHOD AND APPARATUS FOR CACHING NATIVE CODE IN A VIRTUAL  
MACHINE INTERPRETER

INVENTORS: David Wallman  
777 S. Mathilda Ave. #266  
Sunnyvale, California 94087  
A Citizen of Israel

ASSIGNEE: SUN MICROSYSTEMS, INC.  
901 SAN ANTONIO ROAD  
PALO ALTO, CALIFORNIA 94303

A DELAWARE CORPORATION

BEYER WEAVER & THOMAS, LLP  
P.O. Box 130  
Mountain View, CA 94042-0130  
Telephone (650) 961-8300

# METHOD AND APPARATUS FOR CACHING NATIVE CODE IN A VIRTUAL MACHINE INTERPRETER

## 5 BACKGROUND OF THE INVENTION

### 1. Field of Invention

The present invention relates generally to computing systems which use virtual machines. More particularly, the present invention relates to a technique increasing overall virtual machine performance.

### 10 2. Description of the Related Art

The Java™ programming language is an object-based high level programming language developed by Sun Microsystems, Inc. of Palo Alto, California. The Java™ programming language is designed to be portable enough to be executed on a wide  
15 range of computers which include, but are not limited to, small personal computing device and supercomputers.

Programs written in the Java™ programming language, *i.e.*, Java™ programs, are compiled into class files which include virtual machine instructions, *e.g.*,  
20 bytecodes, for a Java™ virtual machine. In essence, a Java™ virtual machine is a software emulator of a “generic” computer. Virtual machine instructions may be run, unmodified, on substantially any computer system that has a virtual machine implementation. As such, virtual machines are both portable and flexible. In addition, other programming languages may be compiled into Java™ virtual machine  
25 instructions and executed by a Java™ virtual machine.

A Java™ virtual machine has an interpreter that is executed as an interpreter program, *i.e.*, software. Conventional interpreters are arranged to decode and to execute the virtual machine instructions of an interpreted program one instruction at a  
30 time during execution. Compilers, which are also associated with a virtual machine, decode source code into native machine instructions prior to execution so that decoding it is not necessary to perform decoding during execution. Conventional interpreters decode each instruction each time the instruction is executed. In other words, conventional interpreters repeatedly decode each instruction each time the

instruction is executed. As such, since an instruction is decoded by a conventional interpreter each time the instruction is encountered, execution of interpreted programs is typically substantially slower than execution of compiled programs since the execution of compile programs may occur on an associated native machine or computer system without necessitating decoding.

As the popularity of small computing devices, *e.g.*, cellular telephones and personal digital assistants, which execute code written in the Java™ programming language increases, the demand for better performing Java™ virtual machines is also increasing. One conventional method for increasing the execution speed and, hence the performance, of Java™-based interpreted programs of virtual machine instructions involves utilizing a just-in-time (JIT) compiler to translate Java™ bytecodes into native code substantially “on-the-fly.” A JIT compiler compiles an entire Java™ function just before the Java™ function is called. However, native code generated by a JIT compiler does not always run faster than code executed by an interpreter. For example, if the interpreter is not spending the majority of its time decoding the Java™ virtual machine instructions, then compiling the instructions with a JIT compiler may not increase the execution speed. In some instances, execution of the virtual machine may even be compromised by utilizing the JIT compiler, *e.g.*, if the overhead associated with compiling the instructions is more than the overhead associated with interpreting the instructions.

In general, at least parts of a JIT compiler are platform specific. As a result, if a virtual machine with a JIT compiler is to be ported to a different platform, at least the platform-specific parts of the JIT compiler, *e.g.*, a code generator, must be rewritten to accommodate the new platform. Rewriting platform-specific parts of a JIT compiler may be a time-consuming task.

Further, a JIT compiler is often associated with a substantial amount of code. That is, JIT compilers often require a relatively significant portion of memory. Even a “simple” JIT compiler such as a JIT compiler which is substantially arranged only to generate code and not to perform extensive optimizations, requires more memory than is typically available in small devices. As will be appreciated by those skilled in

the art, a standard JIT compiler may require memory to store an intermediate representation of compiled code, to store JIT compiler code and data, and to store native code corresponding to entire Java™ methods.

- Therefore, what is desired is a technique for reliably increasing the performance of virtual machines. More particularly, what is needed is a method and an apparatus for increasing the performance of virtual machines in environments with limited memory resources.

## 10 SUMMARY OF THE INVENTION

- The present invention relates to the efficient implementation of a virtual machine in a computing environment. According to one aspect of the present invention, a method for increasing the performance of a virtual machine includes obtaining a program instruction to be executed by the virtual machine, determining when the program instruction is a branch instruction, and determining when a basic block is present in a code cache if the program instruction is a branch instruction. The basic block includes code, and the code cache is associated with the virtual machine. Finally, the method includes executing the code included in the basic block when if the basic block is present. In one embodiment, the program instruction is a bytecode and the code cache is a native code cache.

- In another embodiment, when the basic block is not present in the code cache, the method includes interpreting the program instruction and copying code corresponding to the program instruction into the code cache. In such an embodiment, the method may also include allocating space in the code cache for the code corresponding to the program instruction, in addition to providing the code corresponding to the program instruction with a label.

- According to another aspect of the present invention, a computing system includes a virtual machine which has a code cache and an interpreter. The interpreter is arranged to obtaining a bytecode and to determine when the bytecode is a branch bytecode. The interpreter is also arranged to determine when a basic block is present in the code cache if the bytecode is a branch bytecode. The basic block includes

native code, and the interpreter causes the native code to be executed when it is determined that the basic block is present. In one embodiment, the interpreter is further arranged to interpret the bytecode when it is determined that the bytecode is not a branch bytecode.

5

According to yet another aspect of the present invention, a computer-implemented method for increasing the performance of an interpreter includes identifying a portion of compiled code which corresponds to a block of source code, and copying the portion of compiled code directly into a code cache without first  
10 generating an intermediate representation. In one embodiment, the code cache is a native code cache, and the portion of compiled code is native code.

These and other advantages of the present invention will become apparent upon reading the following detailed descriptions and studying the various figures of  
15 the drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The invention may best be understood by reference to the following  
20 description taken in conjunction with the accompanying drawings in which:

Figures 1a-1e are a diagrammatic representation of a program executed by an interpreter and a stack which stores data associated with the program in accordance with an embodiment of the present invention.

Figure 2a is a diagrammatic representation of instructions of an interpreted  
25 program which are associated with blocks of code in accordance with an embodiment of the present invention.

Figure 2b is a diagrammatic representation of bytecodes, source code, and native code stored in a native code cache in accordance with an embodiment of the present invention.

Figure 2c is a diagrammatic representation of a flow of execution associated with bytecodes, source code, and native instructions in accordance with an  
30 embodiment of the present invention.

Figure 3 is a process flow diagram which illustrates the steps associated with processing a bytecode in accordance with an embodiment of the present invention.

Figure 4 is a diagrammatic representation of an interpreted program, a label table, and a native cache in accordance with an embodiment of the present invention.

5        Figure 5 is a diagrammatic representation of a native block of code with a beginning label and an ending label in accordance with an embodiment of the present invention.

Figure 6 is a diagrammatic representation of a typical, general-purpose computer system suitable for implementing the present invention.

10       Figure 7 is a diagrammatic representation of a virtual machine suitable for implementing the present invention.

#### DETAILED DESCRIPTION OF THE EMBODIMENTS

15       Although just-in-time (JIT) compilers are effective for translating bytecodes, *e.g.*, bytecodes associated with the Java™ programming language developed by Sun Microsystems, Inc., of Palo Alto, California, JIT compilers generally require a substantial amount of memory, and may involve significant overhead. As small computing devices generally have a limited amount of memory, while users all small  
20       computing devices have an expectation of a relatively high level of performance, however, the use of JIT compilers is not always desirable.

By copying native code, *e.g.*, native binary code, corresponding to a bytecode into a cache, and adding interpreter code which labels the code copied into the cache,  
25       the native code may be available on the cache for future use. As will be understood by those skilled in the art, code and data associated with most processor architectures is addressed independently. Hence, when a piece of code in a code segment is copied to a new location, the results of executing the piece of code from within the data segment and from the new location will be the same, provided that no branch  
30       instructions which have references outside the piece of code exist.

In one embodiment, interpreter code which copies native code corresponding to a just-executed bytecode into a native code cache may be added to an interpreted



program. By copying native code into a native code cache, a substantially compiled version of a basic block is then available on the native code cache. Hence, porting is not needed, and the need to recompile native code for different platforms is avoided. Instead, substantially only an interpreter with JIT compilation technology would be recompiled for different platforms.

Generally, an interpreter may execute a bytecode program by repeatedly executing an execute step, an advance step, and a dispatch step. An execute step executes a current bytecode, an advance step advances a bytecode pointer to a subsequent bytecode, and a dispatch step fetches the bytecode identified by the bytecode pointer. The dispatch step also jumps to the piece of native machine code that implements the identified bytecode, e.g., the dispatch operation jumps to an execution step. As will be understood by those skilled in the art, the execution of the execute-advance-dispatch sequence for a bytecode is often referred to as an “interpretation cycle.”

With reference to Figures 1a-e, the execution of a source code statement  $X:=A+B$  after the source code statement is compiled into bytecodes will be described in accordance with an embodiment of the present invention. Instructions 104 associated with the source code statement  $X:=A+B$  may be executed by a virtual machine. As will be understood by those skilled in the art, a virtual machine may be implemented as a stack based machine that includes a stack 110 with frames 114 which are indexed by a stack pointer 118.

Once a first instruction 106a of instructions 104 is interpreted by an interpreter 120, a value of A is pushed onto stack 110 at a location 114a identified by stack pointer 118 which points to a “current” location within stack 110, as shown in Figure 1b. After interpreter 120 executes first instruction 106, interpreter 120 advances to a second instruction 106b, and dispatches second instruction 106b, as shown in Figure 1c. Like first instruction 106a of Figure 1b, second instruction 106b is a load or push instruction. Hence, the execution of second instruction 106b places a value of B onto stack 110 at a location 114g specified by stack pointer 118.

As shown in Figure 1d, once interpreter 120 completes an interpretation cycle associated with second instruction 106d, a third instruction 106c may be executed. In the described embodiment, third instruction 106c is an add operation which adds values A and B, and pushes the result of adding values A and B onto location 114a of stack 110. It should be appreciated that once the sum of values A and B is determined, the result of the sum is stored on stack 110 such that values A and B are displaced from stack 110. Finally, after interpreter 120 completes another interpretation cycle, a fourth instruction 106d is executed, and the result stored in location 114a is removed as shown in Figure 1d. Once the result is removed from location 114a, the result is effectively placed into location 114c which is associated with variable X.

In essence, with respect to Figures 1a-e, after the values of A and B are loaded onto the stack using an interpreted program, the values of A and B may be added and removed from the stack with the result of the summation of A and B being placed on the stack. The result of the summation placed on the stack may then be stored in a location associated with variable X, and the result of the summation may be removed from the stack.

While running an interpreted program, blocks of code associated with instructions may be executed by effectively "jumping to" an appropriate block of code when the associated instruction is executed. Figure 2a is a diagrammatic representation of instructions of an interpreted program which are associated with blocks of code in accordance with an embodiment of the present invention. An interpreted program 202 may include instructions 206 such as a load instruction 206a and an add instruction 206b. While executing interpreted program 202, when an interpreter executes load instruction 206a, interpreter effectively executes a block 210 of code that is associated with load instruction 206a, *i.e.*, block 210a.

As shown, block 210a includes code written in the C programming language that executes a load operation. That is, block 210a contains C code that is associated with interpreted load instruction 206a. The code contained in block 210a may be native C code. Just as code contained in block 210a may be native C code associated

with load instruction 206a, code contained in block 210b may be native C code associated with add instruction 206b. It should be understood that although blocks 210 have been described as including code written in the C programming language, the code associated with blocks 210 may instead be written in substantially any computing language which allows copying of blocks of code segments.

Native code generally includes native machine instructions which are machine instructions that are essentially designed for a specific microprocessor or computer architecture. As will be appreciated by those skilled in the art, a machine instruction is an instruction that directs a computer to perform an operation specified by an operation code and, in some cases, one or more operands.

Figure 2b is a diagrammatic representation of a system which includes a native code cache in accordance with an embodiment of the present invention. As shown, a stream of bytecodes 252 corresponds to source code 256 for an interpreter. Although source code 256 may be written in substantially any computing language, in one embodiment, source code 256 may be written in the C computer language.

Source code 256 generally includes a switch operator which “executes” bytecodes 252 as long as there are bytecodes 252 to obtain. Source code 256 includes cases which correspond to actions performed on specific arguments associated with bytecodes 252. The cases include case blocks 260 which correspond to specific bytecodes 252. As shown, bytecode 252a corresponds to case block 260a, bytecode 252b corresponds to case block 260b, and bytecode 252c corresponds to case block 260c.

Conventionally, case blocks 252 would generally be compiled into compiled code or an intermediate representation, which would then be transformed into instructions, e.g., native code, which would be executed by a processor. In the described embodiment, however, compiled code, or an intermediate representation of source code 256, is not generated. Instead, case blocks 252 may be copied substantially directly into a native code cache 264. It should be understood by those

skilled in the art that case blocks 252 generally include only indirect calls due to addressing issues.

To facilitate the copying of case blocks 252 substantially directly into native code cache 264, labels may be implemented around case blocks 252 to delineate the beginning and end of a specific case block, *e.g.*, case block 252c. By way of example, a beginning label 270 and an ending label 272 may be inserted before case block 252 and at the end of case block 252, respectively. Ending label 272 is placed before a break instruction which typically terminates each case. The use of labels 250, 252 enables code which is to be copied into native code cache 264 to be denoted and, hence, readily identified.

In addition to having boundaries on case blocks 252, or blocks of source code, such that case blocks 252 may be substantially exactly identified, case blocks 252 are arranged to be manipulated at a source code level. The ability to manipulate case blocks 252 at a source code level enables a target to be substantially platform independent.

As shown in Figure 2c, when a stream 280 of bytecodes 282 is interpreted by an interpreter 283 which includes case blocks 284, the case blocks 284 may be copied substantially directly, *e.g.*, without the creation of an intermediate representation, into native code cache 286. When a processor finishes executing a code segment, *e.g.*, a segment including case blocks 284a-c, in native code cache 286, a jump back instruction 288 in native code cache 286 effectively returns control to interpreter 283, which has no control while a processor executes code segments in native code cache 286.

With reference to Figure 3, the processing of a bytecode associated with a program will be described in accordance with an embodiment of the present invention. A process 300 begins when the "next" instruction or bytecode in a program is to be executed. In step 302, a determination is made as to whether there is a bytecode to obtain. If there is no bytecode to obtain, then the process terminates.

Alternatively, when it is determined that there is a bytecode to obtain, then process flow proceeds to step 303 in which the next bytecode to obtain is obtained.

- After the bytecode is obtained, a determination is made in step 304 regarding
- 5 whether the obtained bytecode, *i.e.*, the next bytecode to be executed, is a branch bytecode. A branch bytecode, in one embodiment, may be a bytecode that has a corresponding section or block that is to be run, as described above with respect to Figure 2. If it is determined in step 304 that the bytecode is not a branch bytecode, then the bytecode is a "regular" bytecode. Accordingly, in step 308, the bytecode is
- 10 interpreted. Once the bytecode is interpreted, then process flow returns to step 302 and the determination of whether there are more bytecodes to be obtained.
- Alternatively, if it is determined in step 304 that the bytecode is a branch bytecode, then the implication is that the bytecode has an associated block of code. As such, process flow proceeds to step 312 in which a target is computed. Computing a target
- 15 involves interpreting bytecodes.

- Once a target is computed in step 312, a determination is made in step 316 as to whether a basic block which corresponds to the target is found in a native code cache. In one embodiment, such a determination may be made by searching through
- 20 a "label" table which effectively lists labels associated with basic blocks which are present in the native cache. Specifically, the target may be compared against labels in the label table to determine if a corresponding basic block, *e.g.*, case block 260c of Figure 2b, is present in the native code cache. If it is determined in step 316 that a basic block corresponding to the target is cached in the native code cache, the code
- 25 associated with the basic block is executed in step 320. After the code is executed, process flow returns to step 302 in which a determination is made as to whether there is another bytecode to obtain. Typically, the code associated with the basic block ends with a branch, or jump, instruction to return from the basic block to approximately the point in the program which resulted in the execution of the basic
- 30 block. A basic block is generally a sequence of instructions which start and end at the same place. That is, a basic block is a section of code with no branching and no internal labels. Typically, a basic block ends with a branch instruction.

When it is determined in step 316 that a basic block corresponding to the target is not already in the native code cache, process flow moves from step 316 to step 324 in which a new basic block may be labeled. In the described embodiment, labeling the new basic block also includes indexing the label in the label table, as well as allocating space in the native code cache. Substantially any suitable method may be used to allocate space in the native code cache, as will be appreciated by those skilled in the art. Once space is allocated in the native code cache, then the bytecode is interpreted in step 328. After the bytecode is interpreted, the native code corresponding to the bytecode is copied into the native code cache in step 332.

Specifically, the native code corresponding to the bytecode is copied into the native code cache in the space that was allocated in step 324. Once the native code is copied, process flow returns to step 302 and the determination of whether there is another bytecode to obtain.

In one embodiment, code which allows the copying of the native code to be performed may be inserted within a program. Such code may effectively be code which determines where a block of native code is located, identifies where the block begins and ends, and copies the native code. The code which allows copying may be written in substantially any suitable programming language, as for example, the C programming language.

The copying of native code into a native code cache effectively removes the need for a compiler such as a JIT compiler. That is, no intermediate representation, *e.g.*, no compiled representation, is needed to generate native code, and native code may be obtained substantially immediately. As such, the compilation overhead associated with a compiler such as a JIT compiler may be avoided.

As mentioned above with respect to Figure 3, a label table is used in one embodiment to determine whether a particular basic block of native code is cached in a native cache. Figure 4 is a diagrammatic representation of an interpreted program, a label table, and a native cache in accordance with an embodiment of the present invention. When an instruction or a bytecode 406 in a program 402 is to be executed, it is determined whether there is a basic block 418 associated with bytecode 406. A

label table 410 is used to effectively map bytecodes 406 to blocks 418 stored within native code cache 414. If there is no entry in label table 410 for a particular bytecode 406, and it is appropriate for the particular bytecode 406 to have an entry, then space may be allocated within native code cache 414 for a block 418 that corresponds to the particular bytecode 406, and a label for the particular bytecode 406 may be created.

In order to facilitate identifying basic blocks to be put in a native code cache, labels or identifiers may be placed at the beginning of blocks. Specifically, a block of code pertaining to a load instruction may have a "beginning" label which identifies the beginning of the block of code, as well as an "ending" label which identifies the end of the block of code. By way of example, as shown in Figure 5, a block of native code 502 may be identified by a beginning label 504 and an ending label 506. Identifying where a block of native code begins and ends enables the block to be readily identified and copied into a native code cache.

As described above with respect to Figure 3, a basic block may end with a branch or a jump instruction. Typically, a basic block may be entered substantially only at a start point, and exited only at the branch instruction. In one embodiment, an extended basic block may be used in lieu of a basic block. Typically, an extended basic block is a sequence of instructions that may be entered at more than one start point, but includes only one exit point.

Table 1 below provides a code listing which implements the present invention. It should be understood that the code shown in Table 1 is an example of an embodiment of the present invention. In general, substantially any language which enables blocks of code to be copied may be used to implement the present invention.

---

```

/*****
this is an interpreter implementing a very simple machine:

there are 2 registers x and X
comands with capital letter are applied to X
comands with small letter are applied to x

x? - load ? to x, where ? is a digit
X? - load ? to X, where ? is a digit
a? - add ? to x , where ? is a digit or X
A? - add ? to X, where ? is a digit or x

```

```

m - store in mem value of x - like calculator does
r - restore x from mem
M - store in mem value of X - like calculator does
R - restore X from mem
5 p - print content of x
P - print content of X
l?# - loop command, where ? and # are digits, the machine executes
last ? commands (not including l?#) # times
-- NO NESTED LOOPS --
10 S - print machine state (for debugging)

compile jit and/or interpreter :
gcc Jit.c -g -DJIT=1 -o jit
gcc Jit.c -g -o int

15 cmd Line for inc 1..10 : jit x0 a1 p l29
cmd Line for Fibonacci(10) : jit x0 X1 p P m Ax r P l45
cmd Line for Fibonacci(46) : jit x0 X1 p P m Ax r P l49 m Ax r P l49
m Ax r P l49 m Ax r P l49 m Ax r P l44
20 *****/

#if DBG > 2
#define PRINTF( a, args...) printf( a, ##args )
#define PRINT( a, args...) print( a, ##args )
25 #else
#define PRINTF(a, args...)
#define PRINT(a, args...)
#endif

30 #if DBG > 1
#define PR_STATE() printSt(__LINE__);

#else
35 #define PR_STATE()
#endif

#define PRSTATE() printSt(0);

40 void Pr_state(int lineN ) {
extern int cur, lcntr,x,X,m;
extern char *freePtr;
printf( "LN:%d cur=%d lcntr=%d x=%d\n", lineN ,cur,lcntr,x);
/*printf( "LN:%d cur=%d lcntr=%d x=%d X=%d m=%d freePtr=%x\n",\ **/
45 }

#if JIT

#define JSTART(n) \
50 Strt_## n : \
{ \
PR_STATE();

#define JEND(n) \
55 PR_STATE(); \
} \
End_## n: \
pB =(int) &&Strt_## n; \
pE =(int) &&End_## n; \
60 csz = pE - pB; \
bcopy(pB,freePtr,csz); \

```



```

        freePtr += csz; \
        PR_STATE();

5      #else

        #define JSTART(n) PR_STATE();
        #define JEND(n) PR_STATE();

10     #endif

        /***** */

        char ncache[10000];

15     void* nTable[100];
        char *freePtr= ncache;

        char ret[40];
        int rsize;

20     char curpp[20];
        int curppsize;

        void (*print)(char *, int);
        void (*printSt)(int);

25     void Print(char *s, int x) { printf(" %s = %d\n", s,x ); }

        int nothing = 0;

30     /* they are locals, I put it here for debug */
        int cur;
        int lcntr=0;
        int x, X;
        int m=0;

35     main(int argc, char *argv[]) {

        char *tkn;

40

        int pB, pE;
        int csz;
        int x1;

45     print = Print;
        printSt = Pr_state;

        #if JIT
        /* setup the jit
        build ret_jmp to interpreter loop
        */
        {
50     ___label___ StrtR, EndR;
        void *rt = &&RetLabel;

55     if(nothing){
        StrtR:
            goto *rt;
        EndR:
        }
60

```

```

5      pB =(int) &&StrtR;
      pE =(int) &&EndR;
      x1= pE - pB;
      bcopy(pB,ret,x1);
      rsize=x1;

      }
      #endif

10     /* Interpreter */

      for(cur=1; cur< argc; ) {

15         switch ( tkn[0] ) {

            case 'x' :
JSTART(1)             x = (tkn[1]-'0');
                       cur++;
                       tkn = argv[cur];
20         JEND(1)      break;

            case 'X' :
JSTART(2)             X = (tkn[1]-'0');
                       cur++;
                       tkn = argv[cur];
25         JEND(2)      break;

            case 'a' :
JSTART(3)             if( tkn[1] == 'X' )
                       x += X;
                       else
30                         x += (tkn[1]-'0');
                       cur++;
                       tkn = argv[cur];
35         JEND(3)      break;

            case 'A' :
JSTART(4)             if( tkn[1] == 'x' )
                       X += x;
                       else
40                         X += (tkn[1]-'0');
                       cur++;
                       tkn = argv[cur];
45         JEND(4)      break;

            case 'm' :
JSTART(5)             m = X;
                       cur++;
                       tkn = argv[cur];
50         JEND(5)      break;

            case 'r' :
JSTART(6)             x = m;
                       cur++;
                       tkn = argv[cur];
60         JEND(6)      break;

```

```

JEND(6)
                                break;
                                case 'l' :
JSTART(7)
                                if( lcntr > 0 ) {
                                    lcntr--;
                                    if ( lcntr > 0 )
                                        cur -= tkn[1] - '0';
                                    else
                                        cur++;
                                }
                                else {
                                    cur -= tkn[1] - '0';
                                    lcntr = tkn[2] - '0';
                                }
                                tkn = argv[cur];
JEND(7)

#ifdef JIT
PR_STATE();

                                bcopy(ret, freePtr, rsize);
                                freePtr += rsize;

                                /* return here from native after BB is over
                                */

RetLabel:

                                /* loop do jitting */
                                if( nTable[cur] != 0 )
                                    goto *(nTable[cur]);
                                else { /* new nblock */
                                    nTable[cur] = freePtr;
                                }
#endif

                                break;
                                case 'p' :
JSTART(8)
                                print("x", x);
                                cur++;
                                tkn = argv[cur];
JEND(8)
                                break;
                                case 'P' :
JSTART(9)
                                print("X", X);
                                cur++;
                                tkn = argv[cur];
JEND(9)
                                break;
                                case 'S' :
JSTART(10)
                                PRSTATE();
JEND(10)
                                break;
                                default:
                                    cur++; /* ignore bad command */
                                    tkn = argv[cur];

```

}  
}  
}

5

Table 1.

Figure 6 illustrates a typical, general-purpose computer system suitable for implementing the present invention. The computer system 1030 includes at least one  
10 processor 1032 (also referred to as a central processing unit, or CPU) that is coupled to memory devices including primary storage devices 1036 (typically a read only memory, or ROM) and primary storage devices 1034 (typically a random access memory, or RAM).

15 Computer system 1030 or, more specifically, CPUs 1032, may be arranged to support a virtual machine, as will be appreciated by those skilled in the art. One example of a virtual machine that may be supported on computer system 1030 will be described below with reference to Figure 7. As is well known in the art, ROM acts to transfer data and instructions uni-directionally to the CPUs 1032, while RAM is used  
20 typically to transfer data and instructions in a bi-directional manner. CPUs 1032 may generally include any number of processors. Both primary storage devices 1034, 1036 may include any suitable computer-readable media. A secondary storage medium 1038, which is typically a mass memory device, is also coupled bi-directionally to CPUs 1032 and provides additional data storage capacity. The mass  
25 memory device 1038 is a computer-readable medium that may be used to store programs including computer code, data, and the like. Typically, mass memory device 1038 is a storage medium such as a hard disk or a tape which generally slower than primary storage devices 1034, 1036. Mass memory storage device 1038 may take the form of a magnetic or paper tape reader or some other well-known device. It  
30 will be appreciated that the information retained within the mass memory device 1038, may, in appropriate cases, be incorporated in standard fashion as part of RAM 1036 as virtual memory. A specific primary storage device 1034 such as a CD-ROM may also pass data uni-directionally to the CPUs 1032.

CPU's 1032 are also coupled to one or more input/output devices 1040 that may include, but are not limited to, devices such as video monitors, track balls, mice, keyboards, microphones, touch-sensitive displays, transducer card readers, magnetic or paper tape readers, tablets, styluses, voice or handwriting recognizers, or other well-known input devices such as, of course, other computers. Finally, CPU's 1032 optionally may be coupled to a computer or telecommunications network, *e.g.*, an internet network or an intranet network, using a network connection as shown generally at 1012. With such a network connection, it is contemplated that the CPU's 1032 might receive information from the network, or might output information to the network in the course of performing the above-described method steps. Such information, which is often represented as a sequence of instructions to be executed using CPU's 1032, may be received from and outputted to the network, for example, in the form of a computer data signal embodied in a carrier wave. The above-described devices and materials will be familiar to those of skill in the computer hardware and software arts.

As previously mentioned, a virtual machine may execute on computer system 1030. Figure 7 is a diagrammatic representation of a virtual machine which is supported by computer system 1030 of Figure 6, and is suitable for implementing the present invention. When a computer program, *e.g.*, a computer program written in the Java programming language, is executed, source code 1110 is provided to a compiler 1120 within compile-time environment 1105. Compiler 1120 translates source code 1110 into bytecodes 1130. In general, source code 1110 is translated into bytecodes 1130 at the time source code 1110 is created by a software developer.

Bytecodes 1130 may generally be reproduced, downloaded, or otherwise distributed through a network, *e.g.*, network 1012 of Figure 6, or stored on a storage device such as primary storage 1034 of Figure 6. In the described embodiment, bytecodes 1130 are platform independent. That is, bytecodes 1130 may be executed on substantially any computer system that is running on a suitable virtual machine 1140.

Bytecodes 1130 are provided to a runtime environment 1135 which includes virtual machine 1140. Runtime environment 1135 may generally be executed using a processor or processors such as CPUs 1032 of Figure 6. Virtual machine 1140 includes a compiler 1142, an interpreter 1144, and a runtime system 1146. Bytecodes 1130 may be provided either to compiler 1142 or interpreter 1144.

When bytecodes 1130 are provided to compiler 1142, methods contained in bytecodes 1130 are compiled into machine instructions. In one embodiment, compiler 1142 is a just-in-time compiler which delays the compilation of methods contained in bytecodes 1130 until the methods are about to be executed. When bytecodes 1130 are provided to interpreter 1144, bytecodes 1130 are read into interpreter 1144 one bytecode at a time. Interpreter 1144 then performs the operation defined by each bytecode as each bytecode is read into interpreter 1144. That is, interpreter 1144 “interprets” bytecodes 1130, as will be appreciated by those skilled in the art. In general, interpreter 1144 processes bytecodes 1130 and performs operations associated with bytecodes 1130 substantially continuously.

When a method is invoked by another method, or is invoked from runtime environment 1135, if the method is interpreted, runtime system 1146 may obtain the method from runtime environment 1135 in the form of a sequence of bytecodes 1130, which may be directly executed by interpreter 1144. If, on the other hand, the method which is invoked is a compiled method which has not been compiled, runtime system 1146 also obtains the method from runtime environment 1135 in the form of a sequence of bytecodes 1130, then may go on to activate compiler 1142. Compiler 1142 then generates machine instructions from bytecodes 1130, and the resulting machine-language instructions may be executed directly by CPUs 1032. In general, the machine-language instructions are discarded when virtual machine 1140 terminates. The operation of virtual machines or, more particularly, Java virtual machines, is described in more detail in The Java™ Virtual Machine Specification by Tim Lindholm and Frank Yellin (ISBN 0-201-63452-X), which is incorporated herein by reference.

Although only a few embodiments of the present invention have been described, it should be understood that the present invention may be embodied in many other specific forms without departing from the spirit or the scope of the present invention. By way of example, although the present invention has generally been described in terms of a Java virtual machine, it should be appreciated that the present invention may generally be applied to any suitable virtual machine. Examples of suitable virtual machines include, but are not limited to, a Smalltalk virtual machine.

The basic blocks of the present invention, as well as the blocks of code associated with instructions such as load or push instructions, may generally include instructions written in the C programming language. However, it should be understood that blocks of code such as basic blocks may generally be associated with instructions written in substantially any suitable programming language.

The steps associated with executing a bytecode or an instruction may be varied. Steps may be added, removed, altered, and re-ordered. For example, the process of executing a bytecode may be such that when a jump is made to a native code cache in order to execute a block of native code, it may be possible for a jump not to be made back from the block of native code until after a next instruction is obtained.

It should be understood that the copying of native code into a native code cache without the need to generate an intermediate representation of the code is generally applicable to any interpreter which is coded to have distinguishable blocks of source code that is responsible for interpretation of different input instructions. A suitable interpreter would generally also be coded such that there are effectively no jumps from within one block into another. As described above, such an interpreter may be coded to include a switch instruction which uses case blocks, as described above with respect to Figure 2b. Alternatively, such an interpreter may be implemented with a series of "if" and "else if" statements without departing from the spirit or the scope of the present invention. Therefore, the present examples are to be considered as illustrative and not restrictive, and the invention is not to be limited





What is claimed is:

1. A computer-implemented method for increasing the performance of a virtual machine, the computer-implemented method comprising:

obtaining a program instruction to be executed by the virtual machine;

5 determining when the program instruction is a branch instruction;

determining when a basic block is present in a code cache when it is determined that the program instruction is a branch instruction, the basic block including code, the code cache being associated with the virtual machine; and

10 executing the code included in the basic block when it is determined that the basic block is present.

2. A computer-implemented method as recited in claim 1 wherein when it is determined that the basic block is not present in the code cache, the method further includes:

15 interpreting the program instruction; and

copying code corresponding to the program instruction into the code cache.

3. A computer-implemented method as recited in claim 2 further including:

20 allocating space in the code cache for the code corresponding to the program instruction; and

providing the code corresponding to the program instruction with a label.

4. A computer-implemented method as recited in claim 3 further including placing the label in a table of labels.

5. A computer-implemented method as recited in claim 2 wherein determining when the basic block is present in the code cache includes searching through a table of labels to determine if a target associated with the program instruction has a substantially matching label in the table of labels.

6. A computer-implemented method as recited in claim 2 wherein the program instruction is a bytecode, and wherein the bytecode is executed by an interpreter of the virtual machine.

7. A computer-implemented method as recited in claim 2 wherein the code cache is a native code cache, and the code corresponding to the program instruction is native code.

5

8. A computer-implemented method as recited in claim 1 wherein the program instruction is a bytecode and the code cache is a native code cache.

9. A computer-implemented method as recited in claim 8 further including  
10 interpreting the bytecode when it is determined that the program instruction is not the branch instruction.

10. A computer-implemented method as recited in claim 1 further including:  
computing a target using the program instruction, wherein determining when  
15 the basic block is present in the code cache includes determining if the code cache includes any basic blocks which correspond to the target.

11. A computer program product for increasing the performance of a virtual machine, the computer program product comprising:  
20 computer code for obtaining a program instruction to be executed by the virtual machine;  
computer code for determining when the program instruction is a branch instruction;  
computer code for determining when a basic block is present in a code cache  
25 when it is determined that the program instruction is a branch instruction, the basic block including code, the code cache being associated with the virtual machine;  
computer code for executing the code included in the basic block when it is determined that the basic block is present; and  
a computer-readable medium that stores the computer codes.

30

12. A computer program product as recited in claim 11 further including:  
computer code for interpreting the program instruction when it is determined that the basic block is not present in the code cache; and

computer code for copying code corresponding to the program instruction into the code cache when it is determined that the basic block is not present in the code cache.

- 5 13. A computer program product as recited in claim 12 further including:  
computer code for allocating space in the code cache for the code  
corresponding to the program instruction;  
computer code for providing the code corresponding to the program  
instruction with a label; and  
10 computer code for placing the label in a table of labels.
14. A computer program product as recited in claim 11 wherein the program  
instruction is a bytecode and the code cache is a native code cache.
- 15 15. A computer program product as recited in claim 11 wherein the computer-  
readable medium is one selected from the group consisting of a data signal embodied  
in a carrier wave, a floppy disk, a computer memory, a hard disk, an optical disk, a  
tape drive, and a CD-ROM.
- 20 16. A computing system, the computing system including a virtual machine, the  
virtual machine comprising:  
a code cache; and  
an interpreter, the interpreter being arranged to obtaining a bytecode, the  
interpreter further being arranged to determining when the bytecode is a branch  
25 bytecode and to determine when a basic block is present in the code cache when it is  
determined that the bytecode is a branch bytecode, the basic block including native  
code, wherein the interpreter causes the native code to be executed when it is  
determined that the basic block is present.
- 30 17. A computing system including a virtual machine according to claim 16  
wherein the interpreter is further arranged to interpret the bytecode when it is  
determined that the basic block is not present in the code cache and to copy native  
code corresponding to the bytecode into the code cache.

18. A computing system including a virtual machine according to claim 16 wherein the interpreter is further arranged to interpret the bytecode when it is determined that the bytecode is not a branch bytecode.

5

19. A computer-implemented method for increasing the performance of an interpreter, the method comprising:

identifying a portion of compiled code which corresponds to a block of source code; and

10 copying the portion of compiled code into a code cache.

20. A computer-implemented method as recited in claim 19 wherein the code cache is a native code cache, and the portion of compiled code is native code.

15 21. A computer-implemented method as recited in claim 20 further including: executing the portion of compiled code copied into the native code cache using a processor.

22. A computing system comprising:

20 a processor;

an interpreter; and

a native code cache, the native code cache being associated with the processor, wherein the interpreter includes at least one block of binary code that is arranged to be copied substantially directly into the native code cache.

25

23. A computing system according to claim 22 wherein the at least one block of binary code is associated with a bytecode which is arranged to be interpreted by the interpreter.

30 24. A computing system according to claim 23 wherein the at least one block is arranged to include indirect calls.

## ABSTRACT

Methods and apparatus for efficiently implementing a virtual machine are disclosed. According to one aspect of the present invention, a method for increasing  
5 the performance of a virtual machine includes obtaining a program instruction to be executed by the virtual machine, determining when the program instruction is a branch instruction, and determining when a basic block is present in a code cache if the program instruction is a branch instruction. The basic block includes code, and the code cache is associated with the virtual machine. Finally, the method includes  
10 executing the code included in the basic block when if the basic block is present. In one embodiment, the program instruction is a bytecode and the code cache is a native code cache.




FIGURE 1a






FIGURE 1b






FIGURE 1c




FIGURE 1d




FIGURE 1e






Figure 2a

000100 00011000

BYTECODES

SOURCE CODE

COMPILED  
CODE

NATIVE CODE




FIGURE 2b


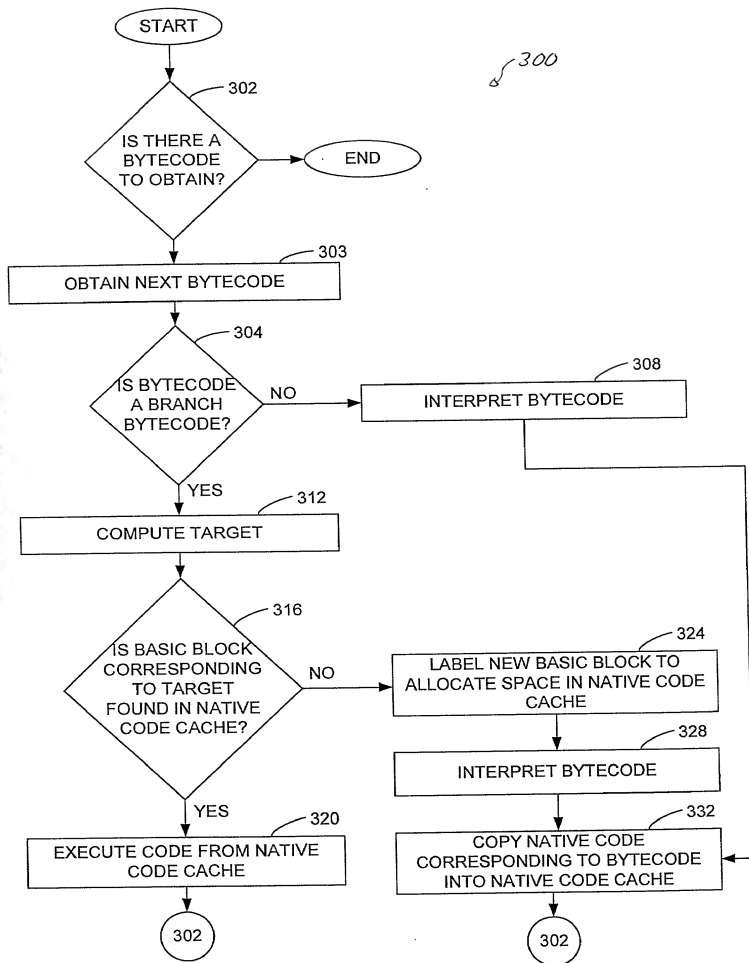


FIGURE 2c



### Figure 3




Figure 4

000100-50041900




Figure 5




Figure 6




Figure 7



# DECLARATION AND POWER OF ATTORNEY FOR ORIGINAL U.S. PATENT APPLICATION

Attorney's Docket No. SUN1P275/P4783

As a below-named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: METHOD AND APPARATUS FOR CACHING NATIVE CODE IN A VIRTUAL MACHINE INTERPRETER the specification of which,

(check one)

1. ☒ is attached hereto.
2. ☐ was filed on \_\_\_\_\_ as  
U.S. Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.
3. ☐ was filed on \_\_\_\_\_ as  
International PCT Application No. \_\_\_\_\_  
and was amended on \_\_\_\_\_.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, CFR § 1.56.

## Prior Foreign Application(s)

I hereby claim foreign priority benefits under Title 35, United States code, § 119(a)-(d) or § 365(b) of any foreign application(s) for patent or inventor's certificate, or § 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application having a filing date before that of the application on which priority is claimed:

\_\_\_\_\_  
(Application No.)                      (Country)                      (Filing Date)

Priority Benefits Claimed?  
Yes \_\_\_ No \_\_\_

\_\_\_\_\_  
(Application No.)                      (Country)                      (Filing Date)

Yes \_\_\_ No \_\_\_

## Provisional Application(s)

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

\_\_\_\_\_  
(Application No.)                      (Filing Date)

\_\_\_\_\_  
(Application No.)                      (Filing Date)

**Prior U.S. Application(s)**

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s), or § 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

_____ (Application No.)	_____ (Filing Date)	_____ (Status - patented, pending, abandoned)
_____ (Application No.)	_____ (Filing Date)	_____ (Status - patented, pending, abandoned)

**Power of Attorney**

And I hereby appoint the law firm of **Beyer Weaver & Thomas, LLP** and all practitioners who are associated with the Customer Number 022434 as my principal attorneys to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

**Direct Correspondence To:**

**Customer Number: 022434**  
**BEYER WEAVER & THOMAS, LLP**  
P.O. Box 130  
Mountain View, CA 94042-0130



**Direct Telephone Calls To:**

**Peggy A. Su (650) 961-8300**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

**Typewritten Full Name of**

**Sole or First Inventor:** David Wallman

**Citizenship:** Israel

**Inventor's signature:** 

**Date of Signature:** 08/14/2000

**Residence:** (City) Sunnyvale

**(State/Country)** CA/US

**Post Office Address:** 777 S. Mathilda Ave. #266, Sunnyvale, CA 94087